



(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 959 416 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
24.11.1999 Bulletin 1999/47

(51) Int. Cl.⁶: G06F 17/30

(21) Application number: 99107926.0

(22) Date of filing: 22.04.1999

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(72) Inventors:
• Leymann, Frank, Dr.
71134 Aldlingen 2 (DE)
• Roller, Dieter, Dipl.-Phys.
71101 Schönaich (DE)

(30) Priority: 08.05.1998 EP 98108425

(74) Representative:
Rach, Werner, Dr.
IBM Deutschland
Informationssysteme GmbH,
Patentwesen und Urheberrecht
70548 Stuttgart (DE)

(71) Applicant:
International Business Machines
Corporation
Armonk, NY 10504 (US)

(54) Using database query technology for message subscriptions in messaging systems

(57) The present invention relates to messaging systems in general and message broker systems in specific. More particularly, the invention relates to means and a method for processing message publication requests and message subscription requests.

The current invention transforms the subscription problem into a database query problem allowing to use query techniques known and optimized for many years. The basic approach of the current invention is to invert the relationship of message publication requests (treated as database data) and message subscription requests (treated as database queries) according the state of the art. Thus the core of the current invention is to treat incoming messages publication requests as database queries and to treat message subscription requests as database queries.




FIG. 3

EP 0 959 416 A2

Description

1 Background of the Invention

1.1 Field of the Invention

[0001] The present invention relates to messaging systems in general and message brokers in specific. More particularly, the invention relates to means and a method for processing message publication requests and message subscription requests.

1.2 Description and Disadvantages of Prior Art

[0002] One source of technology the current invention is based on is the database technology. For a thorough outline of the principles and concepts of database technology refer for instance to R. Elmasri and S.B. Navathe, Fundamentals of database systems, 1989, Benjamin/Cummings. Of specific importance to the current invention is the fundamental concept of a query of database technology. A query is a request to a database system to retrieve from a database a collection of data qualifying under the query; this collection is called the answer set of the query. Typically, the database contains data about real-world entities like persons, orders, appointments etc.. A query is a prescription for selecting subsets from these entities, possibly combine some of them, and present them in a certain form.

[0003] Another source of technology the current invention is based on is the technology messaging systems or message brokers in specific. A messaging system or a message broker is like a hub where messages are streaming in, are transformed, and are streaming out. Messaging systems offer the passing of messages through a store-store-and-forward functionality. Message brokers in addition also allow to reduce the connection costs by supporting the many-to-many communication principle comprising the one "source-message" vs. multiple "sink-message" communication case. Compared to messaging systems a message broker might be augmented by related services like: a message transformer, flow controller, message warehouse, message dictionary and management facility.

[0004] The messages which are streaming into the messaging system/broker are referred to as being published. Messages which are streaming out of the messaging system/broker are referred to as being subscribed. A publication request specifies a message a certain application intends to offer to other interested applications, i.e. to subscribers. A subscription request specifies the subset of all incoming messages a particular application is interested in, and the format in which it has to be presented to the subscriber. An overview on messaging systems and message brokers is given in R. Schulte, Message brokers: A focussed approach to application integration, Gartner Group, Strategic Analysis Report SSA R-401-102, 1996 as well as in "Enter-

prise application integration and information flow", The Yankee Group - Internet Computing Strategies (Boston MA, 1997).

[0005] Thus, a subscription request is similar to what has been described above as a "query" in terms of database technology. And exactly in this sense current state of the art messaging systems and message brokers do treat subscription requests. Expressed in database terms the "messages" are treated as "database data" and the "subscriptions request" are treated as "query" operating on that data filtering out the subscribed messages; i.e. based on its specification messages are filtered, transformed and delivered to its requester. State-of-the-art implementations (for instance the Gryphon approach of IBM Watson Research) use techniques like variants of binary trees to represent subscription requests and process incoming messages accordingly for filtering.

[0006] So far specific implementations for messaging systems and message brokers of these administrations and filtering tasks of message publication and subscription requests have been required. No exploitation of general purpose technologies delivering performant solutions has been possible so far.

[0007] The message broker paradigm is about to gain increased importance in the area of application integration. Its implementation therefore has to be further improved to be able to support the huge diversity of application needs.

1.3 Objective of the Invention

[0008] The invention is based on the objective to improve the processing of message publication requests and message subscription requests within messaging systems in general and message brokers in specific.

2 Summary and Advantages of the Invention

[0009] The objective of the invention is solved by the independent claims.

[0010] The objective of the invention is solved by claim 1. A messaging system is taught comprising first receiving means for receiving a message publication request comprising a message to be published. Said messaging system comprising second receiving means for receiving one or a multitude of message subscription requests comprising a query on a message qualifying under said query. Further a database (301) is comprised. Said second receiving means treats said message subscription request as database data and stores said message subscription request to said database (301). Said first receiving means treats said message publication request as a database query (302) and retrieves from said database a collection (303) of message subscription requests qualifying under said database query.

[0011] The technique proposed by the current inven-

tion makes database technology available to messaging systems. In database environments a query operates on a large amount of data, but in a message broker environment a large amount of queries (corresponding to a message subscription) operate on a single data item (corresponding to a publication of a message). Due to these circumstances database technology could not be exploited in a performant manner by messaging systems. Moreover the current teaching significantly simplifies the implementation of messaging systems as an exploitation of the mature, sophisticated and high performant database technology is possible; no implementation specific to messaging systems is required. Through the transformation of the subscription problem into a database query problem the current invention allows to use query techniques known and optimized for many years. All advantages and benefits provided by database technology are immediately available to messaging systems, like persistency, recovery features, transactional behavior and many more. For a broad spectrum of business problems, no new algorithms must be invented and existing implementations can be reused. All these advantages are of specific value within messages systems required to provide high reliability against data losses and high availability.

[0012] Additional advantages are accomplished by claim 2. According to a further embodiment of the proposed invention said second receiving means transforms a message subscription request into a database data representation.

[0013] By explicitly transforming a message subscription into a specific data representation the efficiency (for instance in terms of the amount of data to be stored or in terms of performance to process a later query for that data etc.) of the current approach can be increased further.

[0014] Additional advantages are accomplished by claim 3. According to a further embodiment of the proposed invention said first receiving means transform said message publication request into a database query.

[0015] Through such a transformation, which may be adjusted to the transformation of the subscription according to claim 2, into a specific query representation the efficiency of the current approach can be increased further.

[0016] Additional advantages are accomplished by claim 4. According to a further embodiment of the proposed invention said first receiving means store said message into said database or into a separate message database.

[0017] Using database technology also for the messages allows to profit from the database technology benefits for the message request too; examples are: the recovery features after system crashes, atomicity, integrity etc.. Moreover through this feature the door to message warehousing is opened.

[0018] Additional advantages are accomplished by

claim 5. According to a further embodiment of the proposed invention the current teaching is applied to a message broker and/or a message warehouse.

[0019] Through its universality the current teaching may be applied in various areas.

[0020] Additional advantages are accomplished by claim 6.

[0021] According to a further embodiment of the proposed invention said database may be stored in main memory and/or on persistent memory.

[0022] Again this demonstrates the flexibility of the teaching; messaging systems, for which the response time behavior is the primary requirement and which will store the database within main memory, as well as messages systems, for which reliability against data losses is the primary requirement and which will store the database on persistent memory, may exploit the proposed technology.

[0023] Additional advantages are accomplished by claim 7. According to a further embodiment of the proposed invention said database may be a relational database or a hierarchical database or a network database or an object-oriented database or a collection of flat files.

[0024] Thus the current teaching is applicable to any type of database. The teaching does not limit the exploiter to a certain database technology; instead other criteria, for instance maturity, availability of the database technology or the type of data to be stored within the database, may decide on the database type to be selected.

[0025] Additional advantages are accomplished by claim 8. According to a further embodiment of the proposed invention the messaging system is comprising sending means for sending said message to subscribers of said collection of message subscription requests.

[0026] Thus the message comprised in the message publication request and serving as the database query can finally be sent to all subscribers of that particular message or that message type.

[0027] The objective of the invention is solved by claim 9.

[0028] A method of processing a message publication request is being taught comprising a message to be published, and one or a multitude of message-subscription-requests comprising a query on a message qualifying under said query. Said method comprises the following steps: 1) a first step of treating said message subscription request as database data and storing said message subscription request to said database (301), and 2) a second step of treating said message publication request as a database query (302) and retrieving from said database a collection (303) of message subscription requests qualifying under said database-query.

[0029] Benefits related to this teaching have been discussed together with claim 1 above.

[0030] Additional advantages are accomplished by

claim 10. According to a further embodiment of the proposed invention said first step is preceded by a subscription transformation step wherein said message subscription request is transformed into a database data representation.

[0031] Benefits related to this teaching have been discussed together with claim 2 above.

[0032] Additional advantages are accomplished by claim 11. According to a further embodiment of the proposed invention said second step is preceded by a message transformation step wherein said message publication request is transformed into a database query.

[0033] Benefits related to this teaching have been discussed together with claim 3 above.

[0034] Additional advantages are accomplished by claim 12. According to a further embodiment of the proposed invention said first step stores said message into said database or into a separate message database.

[0035] Benefits related to this teaching have been discussed together with claim 4 above.

[0036] Additional advantages are accomplished by claim 13. According to a further embodiment of the proposed invention said method is executed by a message broker and/or wherein said database is part of a message warehouse.

[0037] Benefits related to this teaching have been discussed together with claim 5 above.

[0038] Additional advantages are accomplished by claim 14. According to a further embodiment of the proposed invention said method is storing said database in main memory and/or on persistent memory.

[0039] Benefits related to this teaching have been discussed together with claim 6 above.

[0040] Additional advantages are accomplished by claim 15. According to a further embodiment of the proposed invention said database is a relational database or a hierarchical database or a network database or an object-oriented database.

[0041] Benefits related to this teaching have been discussed together with claim 7 above.

[0042] Additional advantages are accomplished by claim 16. According to a further embodiment of the proposed invention said method is further comprising a third step of sending said message to subscribers of said collection of message subscription requests.

[0043] Benefits related to this teaching have been discussed together with claim 8 above.

3 Brief Description of the Drawings

[0044]

Figure 1 is a diagram reflecting the query concept according database technology

Figure 2 visualizes the relationship and the handling of message publication requests and message subscription requests according to

the current state of the art of messaging systems

Figure 3 depicts the relationship and the handling of message publication requests and message subscription requests according to the current invention and applicability of database technology

4 Description of the Preferred Embodiment

[0045] If the current specification is referring to a database this term includes any type of database system, like for example relational databases, hierarchical database, network databases or object-oriented databases; the invention is not limited to a specific database system type. Moreover for the sake of the current invention the data model of the database system is irrelevant. Also, the database system might manage persistent data on disk and/or might manage data in main memory (i.e. a main memory database).

[0046] This teaching does not impose any limitations to the nature of a message or the type of subscription. Moreover the term "message" is to be understood from a logical perspective. The actual communication middleware may transfer the message using the connection-oriented or the connectionless paradigm; the invention is not limited to any of these paradigms.

[0047] The current invention relates to messaging systems as well as message brokers. Therefore if features of the invention are described within this specification in the context of message brokers only the teaching may also be applied to messaging systems and vice versa.

4.1 Introduction

[0048] Of specific importance to the current invention is the fundamental concept of a query of database technology, which is depicted in Figure 1. A query (101) is a request to a database system to retrieve from a database (102) a collection of data qualifying under the query; this collection is called the answer set (103) of the query. Typically, the database contains data about real-world entities like persons, orders, appointments etc.. A query is a prescription for selecting subsets from these entities, possibly combine some of them, and present them in a certain form.

[0049] A messaging system or a message broker is like a hub where messages are streaming in, are transformed, and are streaming out. Messaging systems offer the passing of messages through a store-store-and-forward functionality. Whereas message brokers in addition also allow to reduce the connection costs by supporting the many-to-many communication principle. The messages which are streaming into the messaging system/broker are referred to as being published. Messages which are streaming out of the messaging system/broker are referred to as being subscribed. A publication request specifies a message a certain

application intends to offer to other interested applications, i.e. to subscribers. A subscription request specifies the subset of all incoming messages a particular application is interested in, and the format in which it has to be presented to the subscriber. Message brokers also may support message dictionaries to hold meta-data descriptions of messages for development purposes for instance, and administration and monitoring tools to manage the operation of the broker configuration.

[0050] For example, various stock exchanges might publish stock data periodically, i.e. the stock data is sent to the message broker. Each stock exchange might use a different format to transfer its data. A subscriber has registered to all messages about IBM stock data exceeding \$103 and requests its delivery in XML (extended Markup Language) format.

[0051] Thus, a subscription request is similar to what has been described above as a "query" in terms of database technology. And exactly in this sense current state of the art messaging systems and message brokers do treat subscription requests. A visualization of this state of the art is given in Figure 2. Expressed in database terms the "messages" (202) are treated as "database data" and the "subscriptions request" (201) are treated as "query" operating on that data filtering out the subscribed messages (203); i.e. based on its specification messages are filtered, transformed and delivered to its requester. State-of-the-art implementations (for instance the Gryphon approach of IBM Watson Research) use techniques like variants of binary trees to represent subscription requests and process incoming messages accordingly for filtering.

[0052] Messaging systems and/or message broker support request/reply and/or event-based communication depending on whether the communication is initiated by the a source or a sink application. A sink asks the system for data and the system fulfills the request. In event-based communication communication is source driven. The message is generated by the source and passed to the system for relaying to the sinks. The current invention applies to both communication models.

4.2 Further Analysis of the Publish-Subscribe Problem Area

[0053] The following two problem areas can be identified and are addressed by the current invention.

4.2.1 Many-Queries-Single-Data-Item Characteristic

[0054] Within the area of messaging systems typically many applications have subscribed with a message broker corresponding to many different queries which have to be processed for each incoming message; i.e. message subscriptions requests correspond to database queries. Thus, the situation is the opposite as in data-

base environments: In the latter environments a query operates on a large amount of data, but in a message broker environment a large amount of queries operate on a single data item. Extensive caching of messages can be helpful here, but even if applied to the current situation, the fundamental characteristics, "relatively" small number of messages (database data) versus "relatively" large number of message subscriptions (database queries), is not reversed. Current query processing algorithms within messaging systems do not address this situation, thus, new algorithms must be invented or algorithms from other areas must be applied in this situation.

4.2.2 Subscriptions In Message Warehousing

[0055] In another class of message broker applications the storage of published messages is required resulting in a so-called message warehouse. In this situation, performing the filtering outside of the database as before seems to be unnatural since the database functionality is already part of the overall environment.

[0056] A first step for improvement in this problem area is to store the message in the database before satisfying the subscriptions. Subscriptions will then be transformed into usual database queries and the queries are run against the database to retrieve all corresponding messages. This technique might work fine in certain scenarios: For example, if subscriptions do not have to be considered in real-time the broker can periodically work on the subscriptions, i.e. when a bulk of new messages has been materialized in the database.

4.3 Inverting the Relationship of Message Publication and Subscription Requests

[0057] The current invention transforms the subscription problem into a database query problem allowing to use query techniques known and optimized for many years. Even when messages are not made persistent on disk, query technologies according the current teaching can be applied by considering messages and subscriptions as data in main memory databases.

[0058] The basic approach of the current invention is to invert the relationship of message publication requests (treated as database data) and message subscription requests (treated as database queries). Thus the core of the current invention is to treat incoming message publication requests as database queries and to treat message subscription requests as database data.

[0059] Such an approach is counterintuitive as a message subscription inherently comprises the semantics of an "inquiry", a "filter", "question" and therefore a "query" for a specific type of information, a message. A message on the other hand inherently is of the nature of a data element, as it represents this type of information the subscriber is interested in and which will be sent to

the subscriber.

[0060] The current teaching is depicted in Figure 3. According the current teaching, the messaging system or message broker comprises a database schema which reflects the possible types of subscriptions which can be requested from the message broker. This database schema may depend on the message types published to the broker. Each subscription request received by the broker is transformed into a collection of tuples and stored in the database as database data (301). Without loss of generality one may exploit an underlying relational database being one of the most popular database concepts. Nevertheless the invention may also be based on other data models and database systems. Furthermore, the database must not be persistent, i.e. it may be a main memory database in case no warehousing is required. Various mechanisms exist how to transform queries into tuples. As an example the teaching of F. Leymann, UDH: A universal relation system, Data & Knowledge Engineering 5 (1990), pp. 21 - 38 may be exploited suggesting for the individual predicates a linear representation comprising the search conjunctions and tuple conjunctions. This invention does only assume the existence of one transformation technique, which can be applied, and does not depend on a special one chosen.

[0061] Whenever a message is published, the broker transforms the message into a query (302). This query is performed on the subscription database and its answer set is the collection of subscriptions under which the message qualifies (303). The applications having made the retrieved subscriptions are identified and receive the message in the requested format.

4.4 Value of the Current Approach

[0062] The current invention reduces subscription processing as defined in these messaging environments to query processing known since many years from the field of databases.

[0063] Existing query techniques, implementations, and database systems can be used to perform publish/subscribe in messaging system and message broker environments. For a broad spectrum of business problems, no new algorithms must be invented and existing implementations can be reused. In warehousing environments this teaching is much more efficient than the implementations being done without query processing technology in mind.

Claims

1. A messaging system comprising first receiving means for receiving a message-publication-request comprising a message to be published, and

said messaging system comprising a second receiving means for receiving one or a multi-

tude of message-subscription-requests comprising a query on a message qualifying under said query,

said messaging system being characterized by comprising a database (301), and

wherein said second receiving means treating said message-subscription-request as database data and storing said message-subscription-request to said database (301), and

wherein said first receiving means treating said message-publication-request as a database-query (302) and retrieving from said database a collection (303) of message-subscription-requests qualifying under said database-query.

2. Messaging system according to claim 1,

wherein said second receiving means transforming said message-subscription-request into a database data representation.

3. Messaging system according to any of above claims,

wherein said first receiving means transforming said message-publication-request into a database-query.

4. Messaging system according to any of above claims,

wherein said first receiving means store said message into said database or into a separate message-database.

5. Messaging system according to any of above claims,

wherein said messaging system is a message broker and/or wherein said database is part of a message warehouse.

6. Messaging system according to any of above claims,

wherein said database is stored in main memory and/or on persistent memory.

7. Messaging system according to claim 6,

wherein said database is a relational database or a hierarchical database or a network database or an object-oriented database or a collection of flat files.

8. Messaging system according to any of above claims,
 further comprising sending means for sending said message to subscribers of said collection of message-subscription-requests. 5
9. A method of processing a message-publication-request comprising a message to be published, and
 one or a multitude of message-subscription-requests comprising a query on a message qualifying under said query, 10
 said method comprising the following steps: 15
 a first step of treating said message-subscription-request as database data and storing said message-subscription-request to said database (301), and 20
 a second step of treating said message-publication-request as a database-query (302) and retrieving from said database a collection (303) of message-subscription-requests qualifying under said database-query. 25
10. Method according to claim 1, 30
 wherein said first step is preceded by a subscription-transformation-step wherein said message-subscription-request is transformed into a database data representation. 35
11. Method system according to claim 9 or 10,
 wherein said second step is preceded by a message-transformation-step wherein said message-publication-request is transformed into a database-query. 40
12. Method according to any of claims 9 to 11,
 wherein said first step stores said message into said database or into a separate message-database. 45
13. Method according to any of claim 9 to 12, 50
 wherein said method is executed by a message broker and/or wherein said database is part of a message warehouse.
14. Method according to any of claims 9 to 13, 55
 storing said database in main memory and/or on persistent memory.
15. Method according to claim 14,
 wherein said database is a relational database or a hierarchical database or a network database or an object-oriented database or a collection of flat files.
16. Method according to any of claims 9 to 15,
 further comprising a third step of sending said message to subscribers of said collection of message-subscription-requests.




FIG. 1




FIG. 2




FIG. 3